

Application of the ATT-Meta Metaphor-Understanding Approach to Various Examples in the ATT-Meta Project Databank

Technical Report CSRP-01-02

John A. Barnden

School of Computer Science, The University of Birmingham
Birmingham, B15 2TT, United Kingdom
J.A.Barnden@cs.bham.ac.uk

Abstract

This report provides a measure of relatively objective evaluation of our ATT-Meta approach to conducting some of the reasoning needed in the understanding of metaphorical utterances. The approach is partially implemented in our ATT-Meta system, although the report does not evaluate the system itself. The approach and system are not described in detail here. The utterances of most interest in this report are map-transcending ones: that is, ones that rest on metaphorical views (conceptual metaphors) that are familiar to the understander but transcend the views by using aspects of the source-domain that are not handled by the between-domain mappings the metaphorical views involve. The evaluation of the ATT-Meta approach in this report consists of showing how the approach could extract important target-domain information from mundane, plausibly-map-transcending metaphorical utterances (based on familiar views) that are listed in the author's databank of real-discourse examples of metaphor for mental states and processes. The report shows that the approach successfully produces what the author claims to be the main pieces of information conveyed by the examples. The report also vindicates the ATT-Meta stance that metaphorical mappings should by default not be extended to cope with aspects of utterances that those mappings do not deal with: rather, within-source-domain reasoning should be used to link those aspects to source-domain aspects that those mappings do deal with.

1 Introduction

The aim of this report is to provide a measure of objective evaluation of the “ATT-Meta” approach to conducting some of the reasoning needed in the understanding of metaphorical utterances. This approach is partially implemented in the ATT-Meta system. The approach and system will not be described in detail here—the reader is referred to other papers (Barnden *et al.*, 1996; Barnden, 1998a,b; Barnden & Lee, 1999, 2001a). The present report is intended as an adjunct to another technical report (Barnden & Lee, 2001a), and although the present one is self-contained it is best read after absorbing the less technical portions of the other report.

The metaphorical utterances of most interest in the ATT-Meta project are ones I call *map-transcending*: they rest on metaphorical views (conceptual metaphors) that are familiar to the understander, but transcend them by using source-domain concepts not handled by the between-domain mappings the metaphorical views involve. A possible example is “McEnroe starved Connors to death” when used as a description of a tennis match. A metaphorical view of SPORT AS COMBAT might be familiar to the understander, and that view might map combat-domain concepts such as killing to target-domain concepts such as defeating. However, the existing mapping might not contain any mapping relationship that maps the source-domain concept of starving someone. (I regard a mapping to be a set of mapping relationships, each of which maps an aspect of the source domain to an aspect of the target domain.)

For map-transcending utterances based on familiar metaphorical views, the ATT-Meta approach advocates possibly-extensive inferencing in the terms of the source domain of the metaphorical view, while avoiding as far as possible the extension of the mappings to deal with the map-transcending source-domain aspects of the utterance. Map-extension is regarded as an exceptional action undertaken only when special needs of discourse understanding force it to take place.

The source-domain inferencing is done in an attempt to derive, from the source-domain meaning of the utterance, information that *can* be converted by the existing source-to-target mapping into information about the target domain. The source-domain inferencing is done within a special, protected computational environment I call a *pretence cocoon*. See Figure 1. The cocoon can be thought of heuristically as the belief space of a hypothetical agent who believes that the source-domain meaning of the utterance is true. This approach will become clearer in the examples and diagrams below. In the above starving example, the source-domain meaning is that McEnroe starved Connors to death in the real, biological sense. The inferencing within the cocoon can then establish that McEnroe killed Connors. The existing mapping can then be applied to derive, in the target domain and outside the cocoon, the hypothesis that McEnroe defeated Connors in the match. Other inferences could be made within the cocoon on the basis of the detailed form of killing that starving someone constitutes, and some resulting inferences might be mapped over to the target domain by additional existing mapping relationships.

There is a great deal of understander-relativity in the ATT-Meta approach to metaphor understanding. For one thing, the question of whether a metaphorical view is familiar or not is an understander-relative issue. However, unless otherwise indicated I assume that the particular views mentioned in this report are familiar to ordinary English speakers. The assumption is based partly on the high frequency with which I have observed some of the views being used in mundane discourse, partly on intuition, partly on claims of other metaphor theorists such as Lakoff as to what views are conventional (and as illustrated in compendia such as the Master Metaphor List—Lakoff, 1994), and partly by looking in dictionaries for the phraseology used in utterances.

There is further understander-relativity in the question of what particular mapping relationships between source and target are included in viewing something A as something B. For example, different understanders

may differ in detail on what a metaphorical view of IDEAS AS PLANTS amounts to. I have tried to include in the view only those source/target mapping relationships that seem to be needed in many examples I have encountered.

I also assume that stock English metaphorical phraseology (i.e., phraseology that is frequently occurring and quite fixed, such as “see the problem”) use metaphorical views that are familiar to typical English understanders, and that the source-domain concepts used in the phraseology are mapped by the mappings, so that the phraseology does not count as map-transcending.

Additionally, I will only consider examples in which the metaphoricity is plausibly not *sidelined* for an ordinary English speaker. For example, an understander might have a sports-domain lexicon sense for “kill,” so that he/she/it could use that sense in understanding the sentence “McEnroe killed Connors” in the tennis context. This would circumvent the need to proceed via a sense of kill in the biological combat domain and a mapping from that domain to the domain of sport. Clearly, whether the metaphoricity of an utterance is sidelined or not is again an understander-relative matter.¹

The evaluation of the ATT-Meta approach in this report consists of showing how the approach could extract important target-domain information from (simplified versions of) plausibly-non-sidelined, plausibly-map-transcending metaphorical utterances (based on familiar views) that are listed in a databank of real-discourse examples of metaphor for mental states and processes (Barnden, n.d.). Another constraint on my selection of examples was that I wanted them to come from mundane discourse, or at least to be plausible candidates for mundane discourse. By mundane discourse I mean discourse in conversations, popular magazine articles, popular novels, news articles, popular science texts and other factual texts for wide consumption by the public.

The ATT-Meta approach is not as it stands a specific, complete model of metaphor understanding. It is parametrized by whatever views are familiar to a particular understander, what mapping relationships are included in each familiar view, and what direct, lexicon senses the understander has for words. Thus, it is important to understand that the purpose of this report is not to establish that any particular metaphorical views are familiar to typical understanders, that any particular mapping relationships appear in particular views, or that the metaphoricity in any particular example is not sidelined for typical understanders. The purpose is to show that, given particular decisions about what is familiar, what is mapped and what is not sidelined, the ATT-Meta approach can succeed in drawing useful information about the target domain from a metaphorical utterance; and also to show that very little has to be included in the mappings in order to get great ability to draw out such information. The less sidelining we assume and the fewer mapping relationships we include in a metaphorical view, the more the power of the approach is demonstrated.

I believe the discussion of the individual examples below covers the main informational contributions that the examples make concerning the target domains of the utterances. In most cases I doubt that the suggested informational contributions are controversial.

The plan of the rest of the report is as follows. Section 2 summarizes an important additional feature of the treatment of source/target mappings in the ATT-Meta approach. Section 3 explains why certain metaphorical views in the databank are not covered in the present report. Section 4 applies ATT-Meta to one databank example for each view covered. Section 5 summarizes the performance of the approach on the examples.

The implemented ATT-Meta *system* has not been applied to the examples in this report. However, Lee & Barnden (2001) evaluate the system itself, as well as the approach, by applying it to real-discourse utterances based on most of the metaphorical views about mental state/processes that are listed in the Master

¹Observe that we can still regard “McEnroe killed Connors” as being metaphorical for the understander as long as the biological-combat-domain route is in principle *available* to the understander.

Metaphor List (Lakoff, 1994). See also Barnden & Lee (2001a,b) for detailed descriptions of how the implemented system itself works on some simplified real-discourse examples from the databank. Barnden (2001) evaluates the approach (but again not the system) by applying it to selected examples from Goatly's (1997) comprehensive, detailed survey of types of metaphorical language. Although the examples in these various papers, except the last, concentrate on metaphor of mind, the approach is in no way limited to such metaphor.

2 View-Neutral Mapping Adjuncts

Barnden & Lee (2001a) specify various *view-neutral mapping adjuncts* (VNMA). These provide additional mapping relationships that apply, by default, whatever particular metaphorical view is at hand. They are parasitic on other things that are mapped, in the sense that they are all of the form: if such-and-such things are mapped from the pretence into reality space, then also such-and-such things are mapped.

One of the simplest VNMA is the *Negation* VNMA, which states that if a within-pretence proposition P is mapped into reality space to become a proposition R there, then the negation of P is also mapped, and is mapped to the negation of R. For example, if the proposition that McEnroe killed Connors is mapped to the proposition that he defeated Connors, then, by virtue of the VNMA, the view also maps the proposition that McEnroe did *not* kill Connors to the proposition that he did *not* defeat Connors. Relatedly, the *Qualitative Degree* VNMA maps identically the qualitative degree to which properties and relationships hold.

The *Causation/Ability* VNMA identically maps properties and relationships such as causation, enablement, disablement, ability, assistance, hindrance and easiness. For example, consider the within-pretence proposition that McEnroe's combat actions *enabled* McEnroe to kill Connors *easily*. If the metaphorical view maps combat actions to tennis actions and killing to defeating, then this proposition is mapped to the proposition in reality-space that McEnroe's tennis actions *enabled* McEnroe to defeat Connors *easily*.

The *Change* VNMA maps a change event that applies to some within-pretence state to a change event applying to the corresponding reality-space state if any. The *Event-Shape* VNMA identically maps aspectual features of events/situations/processes in the pretence, such as whether they have a start or end, or are intermittent, to corresponding events/situations/processes if any in the reality space. The *Time-Order* VNMA identically maps the time order of events in the pretence. The *Duration* VNMA identically maps qualitative durations of events.

The *Mental/Emotional States* VNMA maps the mental/emotional states of within-pretence agents over to mental/emotional states of corresponding reality-space agents, if any. The mental attitude (e.g. belief) or the quality of the emotion (e.g. anger) is preserved, but the propositional content, if any, may be modified by other mapping relationships. The *Value-Judgment* VNMA identically maps value-judgments (e.g. by the understander) about within-pretence hypotheses.

The *Function* VNMA identically maps the functions (i.e., purposes served) of entities.

Further VNMA and more refined descriptions are given in Barnden & Lee (2001a). Our VNMA are inspired in part by, but are different from, postulates about mapping put forward by other authors, as detailed in Barnden & Lee (2001a).

3 Choice of Metaphorical Views and Examples

Barnden & Lee (2001a,b) give detailed applications of the ATT-Meta approach, and indeed the ATT-Meta system itself, to examples that use the metaphorical views of MIND AS PHYSICAL SPACE, IDEAS (OR EMOTIONS) AS PHYSICAL OBJECTS, and MIND PARTS AS PERSONS (OR OTHER ANIMATE BEINGS). These are three of the main metaphorical views in the databank. I therefore omit them from the present report.

Also, Barnden & Lee (2001a,b) address the very commonly used view IDEAS AS INTERNAL UTTERANCES to some extent. I omit further discussion of this view in the present report because there are complexities of analysis (Barnden, 1997) that deserve a lengthier treatment than is appropriate here.

I omit discussion of various other views, as follows:

- MIND AS WORLD-DEFINER:
again, because of complexities of analysis
- EMOTIONAL-MIND AS PHYSICAL SPACE:
because the considerations are similar to those for MIND AS PHYSICAL SPACE
- IDEAS AS EXTERNAL UTTERANCES:
because the databank contains few examples of it and I am not confident that it would be familiar to understanders
- IDEAS AS SHARABLE COMMODITIES:
because the databank contains only two examples of it and they do not exhibit interesting map-transcendence
- COGNIZING AS INTERACTING WITH NON-MENTAL ENTITIES:
because the examples in the databank merely use stock phraseology and are therefore not map-transcending.

I also omit the examples in the Miscellaneous Metaphors section of the databank. That leaves the remaining metaphorical views for discussion:

IDEAS AS POSSESSIONS

IDEAS AS PERSONS OR OTHER ANIMATE BEINGS

IDEAS AS EXTERNAL ENTITIES

COGNIZING AS SEEING

MIND AS ANIMATE BEING OR LIVING BODY

MIND AS [INANIMATE] PHYSICAL OBJECT

We will take these in the order listed.

4 Treatment of the Chosen Examples

For each view I have chosen an example of intermediate richness, complexity and difficulty. The simplest examples in the databank for a given view are typically just uses of stock phraseology, and I assume that the utterances are therefore not map-transcending and do not interestingly exercise the ATT-Meta approach. The most complex examples are typically very elaborate ones in refined literary works, whereas for the purposes of this report I am more interested in mundane discourse.

For each example there is a Figure at the end of the report depicting the computational processing sketched in the text.

4.1 Ideas as Possessions

I've inherited his thoughts and ideas. But I've got some [thoughts/ideas] of my own.

[This was said by the new manager of a football club, and "his" refers to the previous manager.]²

I assume that possession of an idea maps to both entertainment of the idea and (when the idea is a proposition) to believing it. The latter could be defeated in a particular case by additional knowledge even if the former is not, but I will assume here that it is not defeated.

See Figure 2. The understander can infer in the pretence cocoon, using simple source-domain knowledge about inheritance, that the new manager has gained possession of the (thoughts and) ideas. But *gain* of possession maps to *starting* to entertain and normally starting to believe the idea, because of the Event-Shape VNMA. Therefore, the approach can produce what is clearly the main informational contribution of the first sentence in the example.

Prototypically, when one inherits something from someone, the latter person has died. Being dead implies being defunct in one's previous relevant role (that of parent, say), and therefore by the Function and Negation VNMAs, the system can infer in our example that the old manager is now defunct as regards his previous relevant role (that of being a manager of the club in question).

There is a further informational contribution that should be discussed. Suppose that in the source-domain of possession it is the case that if X inherits Y from Z then, usually, Z believed that this would happen or intended it to happen. Then, the ATT-Meta approach could defeasibly infer in the pretence cocoon that the old football manager believed that the new manager would inherit his ideas, or intended him to do so. Therefore, by the Mental States VNMA, it follows with some degree of uncertainty that the old manager believed or intended that the new one would entertain/believe his (the old manager's) ideas. It strikes me that this would be a reasonable inference. It is, of course, defeasible.

Finally, the word "inherit" could be claimed to suggest that the new manager had some special relationship to the old one, as a metaphorical "son" perhaps. The new one may have learned a lot from or have been trained by the old one. The inferring of a special relationship of this sort would not arise from the IDEAS AS POSSESSIONS mapping together with VNMAs. However, it would arise if the understander also possessed a metaphorical view of PUPIL AS CLOSE DEPENDANT (using a broad sense of pupil). If it were inferred in the pretence cocoon that the new manager was a close dependant of the old, then it could be inferred that in reality space he was the old one's pupil.

²BBC Radio 4: interview during "Today" news programme, 2 March 1995.

4.2 Ideas as Persons or Other Animate Beings

(2) But [Sara] couldn't look at the whole [mental] image—just jigsaw pieces of it. She saw their mouths, hungry and afraid, but then her thoughts panicked and ran away. [“Their” refers to Sara and a friend, not to the jigsaw pieces.]³

In the second sentence, the clause starting at “but then” uses a familiar view of IDEAS AS PERSONS OR OTHER ANIMATE BEINGS. (The part before this clause is a very interesting example of metaphor, but too complex to comment on here.) I assume that an aspect of IDEAS AS PERSONS OR ... is that physical closeness of the idea-*qua*-person to the agent maps to the agent entertaining the idea.

See Figure 3. The “ran away” implies, by source-domain reasoning, that there was closeness but it ended. So, by the just-mentioned mapping together with the Event-Shape VNMA, the person was entertaining the idea but then ceased to. The rapidity of the running event itself implies rapidity of the closeness-cessation in the source domain, and this rapidity would also map to rapidity of the cessation of mental entertaining, by the Duration VNMA.

The question of the panic is more subtle and difficult. It is not clear whether it is appropriate for the understander to infer that the agent herself panicked. If it is, our approach provides no immediate basis for the inference. This is because the thoughts do not themselves map to any agent, preventing the Mental States VNMA coming into play and allowing the panic of the thoughts to be mapped. It is possible that the role of the panicking is actually purely to stress, within the pretence, the urgency and rapidity of the running-away. This extra rapidity would be carried over by the Duration VNMA application already mentioned. Urgency is a value judgment, and could therefore be mapped by the Value-Judgment VNMA.

I would also claim that “panicked and ran away” could be treated as a unitary event (albeit with two sub-events), corresponding to some single unanalyzed mental event in the target domain, given that panicking and running away commonly go together in ordinary experience. There is no need to postulate two different events in the target domain, corresponding to two events in the source domain. However, the issue takes us into the different area of how events implied by natural language utterances are to be individuated, a difficulty that is not peculiar to metaphorical concerns.

4.3 Ideas as External Entities

In the previous subsection, we implicitly treated the ideas as animate beings external to the agent and running away from her (though a more complex analysis would have been that they were running in a mind-*qua*-physical-space away from the person's conscious self). Therefore we were treating the example as using also the more general metaphorical view of IDEAS AS EXTERNAL ENTITIES. The external entities in this view are usually but not necessarily physical objects. The mapping relationship used, in the previous subsection, between physical closeness and mental entertainment of an idea is inherited from this view. An inanimate example of the more general view is:

I set the notion aside, but I had a feeling it was going to stick to me with a certain burrlike tenacity.⁴

See Figure 4. A physical object that is sticking to a person is as a result close to the person, by default, and hinders being removed. The burr-like tenacity serves to emphasize this hindrance and the longevity of the

³Patti Davis, *Bondage*, p.128. New York: Pocket Books (Simon & Schuster), 1994. Popular novel.

⁴Sue Grafton, *I is for Innocent*, p.182. London: Penguin Books, 1993. Popular detective novel.

sticking. The longevity and resistance could be mapped by the Duration and Causation VNMA. But since the sticking is actually reported in the future tense and also within “a feeling” of the person, we also need to exploit the Mental States VNMA. Then we get the informational contribution that the person has a feeling that she is going to be consciously considering the notion for a long time and that the notion will hinder attempts to stop this entertaining.

4.4 Cognizing as Seeing

It was difficult to concentrate. Someone had turned the dials up and the whole room was writhing with clouds of steam, swirling in [Matthew’s] eyes, his mind.⁵

Because steam is said to be swirling in Matthew’s mind, his mind is being viewed by means of MIND AS PHYSICAL SPACE. It is common for this view to be combined with a form of COGNIZING AS SEEING, where what is seen is mental images inside the mind-space. In this form of COGNIZING AS SEEING the seeing is mapped to conscious consideration of the situations depicted by the images. The “seeing” within the agent’s mind is (I claim) done by the agent’s conscious self. This imaginary person-like entity is claimed by Barnden & Lee (2001a) to be implicated in many uses of MIND AS PHYSICAL SPACE, whether or not they also involve seeing within the space.

For the present example, second sentence, the understander can infer within the pretence cocoon that the steam prevents the conscious self being able to see properly. See Figure 5. Thus, by the Causation VNMA, the understander can obtain the informational contribution that the steam prevents proper conscious consideration (of the states of affairs at issue in the surrounding context).

Notice also that this informational contribution is an elaboration of the previous statement in the discourse that the agent found it difficult to concentrate. Since elaboration is a common discourse-coherence relation (Hobbs, 1985), we have a good example here of how one sentence (the first) could suggest a reasoning goal for guiding the production of informational contributions from another (the second). This is in line with the discourse-query-directed aspect of the ATT-Meta approach as expressed in Barnden & Lee (2001a).

4.5 Mind as Animate Being or Living Body

[Rose] thought, [Gerard] is sick, he is sick, he is *poisoned* by those thoughts, by those *terrible thoughts*.⁶

This uses the familiar view of MIND AS ANIMATE BEING OR LIVING BODY. If we assume that this maps bodily functioning to mental functioning, then the Causation, Change and Value-Judgment VNMA will map a causation of undesirable modification of bodily functioning to causation of undesirable modification of mental functioning. See Figure 6. The source-domain reasoning consists of inferring causation, by the terrible thoughts, of undesirable modification of the bodily functioning from the poisoning. Finally, the sentence is about Rose’s thought about Gerard’s state, not directly about his state. The Mental States VNMA serves to enable the wrapping of everything within the pretence cocoon within a “Rose thought” layer to be carried over to reality-space.

⁵Wendy Perriam, *Born of Woman*, p.129. London: Flamingo (Harper Collins), 1993. Popular/literary novel.

⁶Iris Murdoch, *The Book and the Brotherhood*, p.483. London: Penguin Books, 1988. Italics in original. Literary novel, but passage could be in mundane text by my estimation.

4.6 Mind as Physical Object—Other Cases

It was only momentum that kept Alice’s gears of thought and analysis grinding past any practical application.⁷

This uses the familiar view of MIND AS MACHINE, which is a special case of MIND AS PHYSICAL OBJECT. I assume that one main informational contribution, arising from “momentum,” is that Alice’s thinking activity up to the relevant point in time caused her to continue thinking (without practical application).

See Figure 7. Source-domain reasoning leads to the conclusion that Alice’s mind-machine kept on going but against the resistance of the gear-grinding. The cause of this was the “momentum.” This implies that the machine’s running up to this point is the cause of the continued running. I assume that, under the metaphorical view, running of the machine maps to mental functioning. Then with the aid of the Event-Shape and Causation VNMA’s we get the mentioned informational contribution.

Also, the gear-grinding suggests in the source domain that the machine is tending to stop. This tendency maps to a tendency of Alice’s mind to stop functioning, because of the Event-Shape VNMA.

5 Conclusion

The report, in combination with others (Barnden & Lee, 2001a,b), has shown that the ATT-Meta approach can derive, at least, what are arguably the most important informational contributions from real-discourse metaphor examples based on the most commonly-used metaphorical views represented in the databank (Barnden, n.d.).

The evaluation was restricted to mundane (or at least potentially mundane), non-sidelined, map-transcending utterances based on familiar metaphorical views, as these utterances are the main target of the ATT-Meta approach. Clearly, as I have treated only one example per view in this report, and as I do not claim that the databank is statistically representative in any case, the evaluation is limited and merely suggestive, but the obvious wide diversity of content of the examples provides a measure of assurance that the approach is a promising one in general.

In particular, the report gives suggestive grounds for supposing that the ATT-Meta approach’s default of trying not to do map-extension is on the right track. The information extracted from the example utterances did not require any map-extension, even though the mapping relationships that are assumed to form part of the metaphorical views involved in the utterances are extremely limited in number.

The report did not need to go beyond the particular view-neutral mapping adjuncts suggested in Barnden & Lee (2001a). In fact, the present report did not need the whole set included there on theoretical grounds.

6 Acknowledgments

The research is being supported by grant GR/M64208 from the Engineering and Physical Sciences Research Council (UK).

⁷Jane Smiley, *Duplicate Keys*, p.206. London: Flamingo (Harper Collins), 1996. Popular novel.

7 References

- Barnden, J.A. (1997). Consciousness and common-sense metaphors of mind. In S. O’Nuallain, P. McKeivitt & E. Mac Aogain (Eds), *Two Sciences of Mind: Readings in Cognitive Science And Consciousness*, pp.311–340. Amsterdam/Philadelphia: John Benjamins.
- Barnden, J.A. (1998a). An AI system for metaphorical reasoning about mental states in discourse. In Koenig, J-P. (Ed.), *Discourse and Cognition: Bridging the Gap*, pp.167–188. Stanford, CA: CSLI.
- Barnden, J.A. (1998b). Combining uncertain belief reasoning and uncertain metaphor-based reasoning. In *Procs. Twentieth Annual Meeting of the Cognitive Science Society*, pp.114–119. Mahwah, N.J.: Lawrence Erlbaum Associates.
- Barnden, J.A. (2001). Application of the ATT-Meta metaphor-understanding approach to selected examples from Goatly. Technical Report CSRP–01–01, School of Computer Science, The University of Birmingham, U.K.
- Barnden, J.A. (n.d.). ATT-Meta Project Databank.
<http://www.cs.bham.ac.uk/~jab/ATT-Meta/Databank>.
- Barnden, J.A. (in press). The utility of reversed transfers in metaphor. To appear in *Procs. Twenty-Third Annual Meeting of the Cognitive Science Society* (Edinburgh, August 2001). Lawrence Erlbaum.
- Barnden, J.A., Helmreich, S., Iverson, E. & Stein, G.C. (1996). Artificial intelligence and metaphors of mind: within-vehicle reasoning and its benefits. *Metaphor and Symbolic Activity*, 11(2), pp.101–123.
- Barnden, J.A. & Lee, M.G. (1999). An implemented context system that combines belief reasoning, metaphor-based reasoning and uncertainty handling. In P. Bouquet, P. Brezillon & L. Serafini (Eds), *Second International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT’99), Lecture Notes in Artificial Intelligence, 1688*, pp.28–41, Springer.
- Barnden, J.A. & Lee, M.G. (2001a). Understanding open-ended usages of familiar conceptual metaphors: An approach and artificial intelligence system. Technical Report CSRP–01–05, School of Computer Science, The University of Birmingham, U.K.
- Barnden, J.A. & Lee, M.G. (2001b). Application of the ATT-Meta metaphor-understanding system to an example of MIND PARTS AS PERSONS. Technical Report CSRP–01–09, School of Computer Science, The University of Birmingham, U.K.
- Goatly, A. (1997). *The language of metaphors*. London and New York: Routledge.
- Hobbs, J.R. (1985). On the coherence and structure of discourse. Report No. CSLI–85–37, Center for the Study of Language and Information, Stanford University, Calif., USA.
- Lakoff, G. (1994). The Master Metaphor List. <http://cogsci.berkeley.edu>. University of California, Berkeley.
- Lee, M.G. & Barnden, J.A. (2001). Mental metaphors from the Master Metaphor List: Empirical examples and the application of the ATT-Meta system. Technical Report CSRP–01–03, School of Computer Science, The University of Birmingham, U.K.




Figure 1: **General nature of reasoning in the ATT-Meta approach.** The bold box shows the pretence cocoon. Bold arrows show the action of mapping relationships, between source-domain information in the pretence cocoon and target-domain information in reality. (Target-to-source mapping actions are allowed in the approach, for reasons explained in Barnden, in press. An example occurs in Barnden (2001: Figure 10).) Other arrows, apart from the one at the top, show reasoning actions within reality or within the pretence.

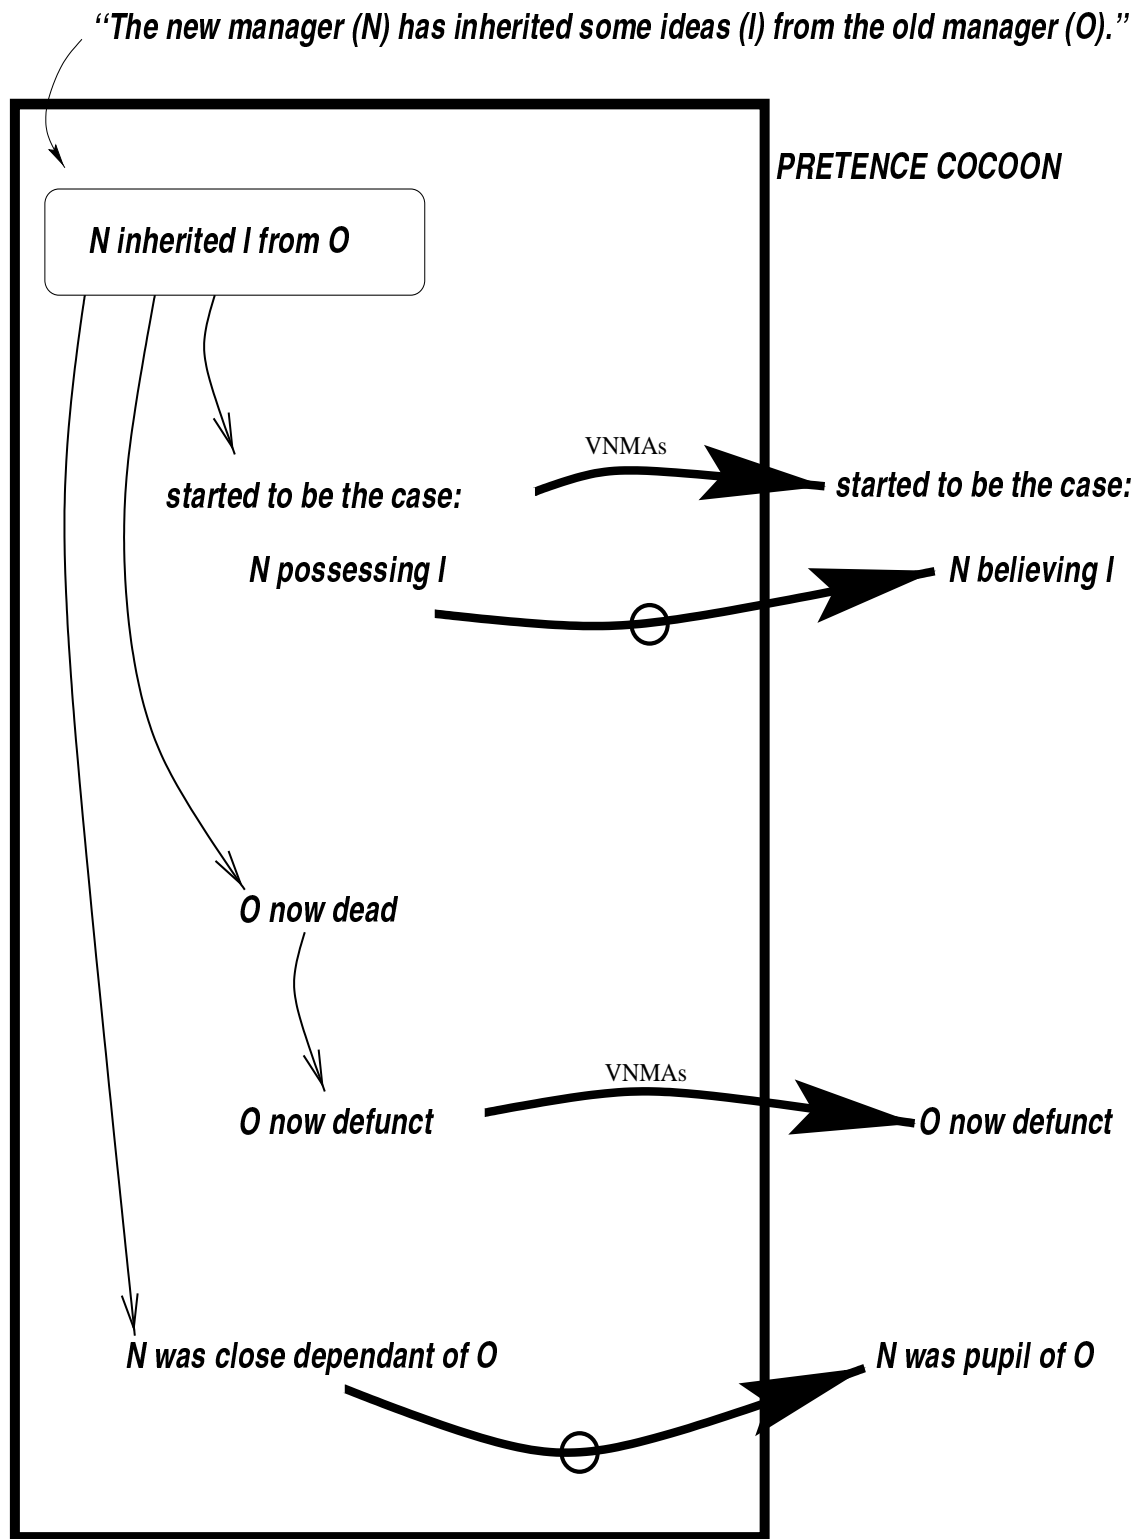


Figure 2: **Ideas as Possessions.** *The Football Inheritance.* In this and later figures, a mapping arrow labelled VNMA or VNMA's shows the action of one or more view-neutral mapping adjuncts. A mapping arrow marked with a circle shows the action of a mapping relationship specific to a particular metaphorical view, here IDEAS AS POSSESSIONS or PUPIL AS CLOSE DEPENDANT. The source-based meaning of the utterance is shown in a small box near the top of the diagram. The propositions within the cocoon box and to its right are English glosses of expressions in some internal representation scheme used by the understander. Tense in the English glosses is actually handled by the Time-Order VNMA (together with identity mapping of the present time point). Therefore, the mapping arrows marked by a circle as view-specific often represent in actuality the combined operation of a view-specific mapping relationship and the Time-Order VNMA. In the figures only a selection of the possible propositions and inferential links are shown.

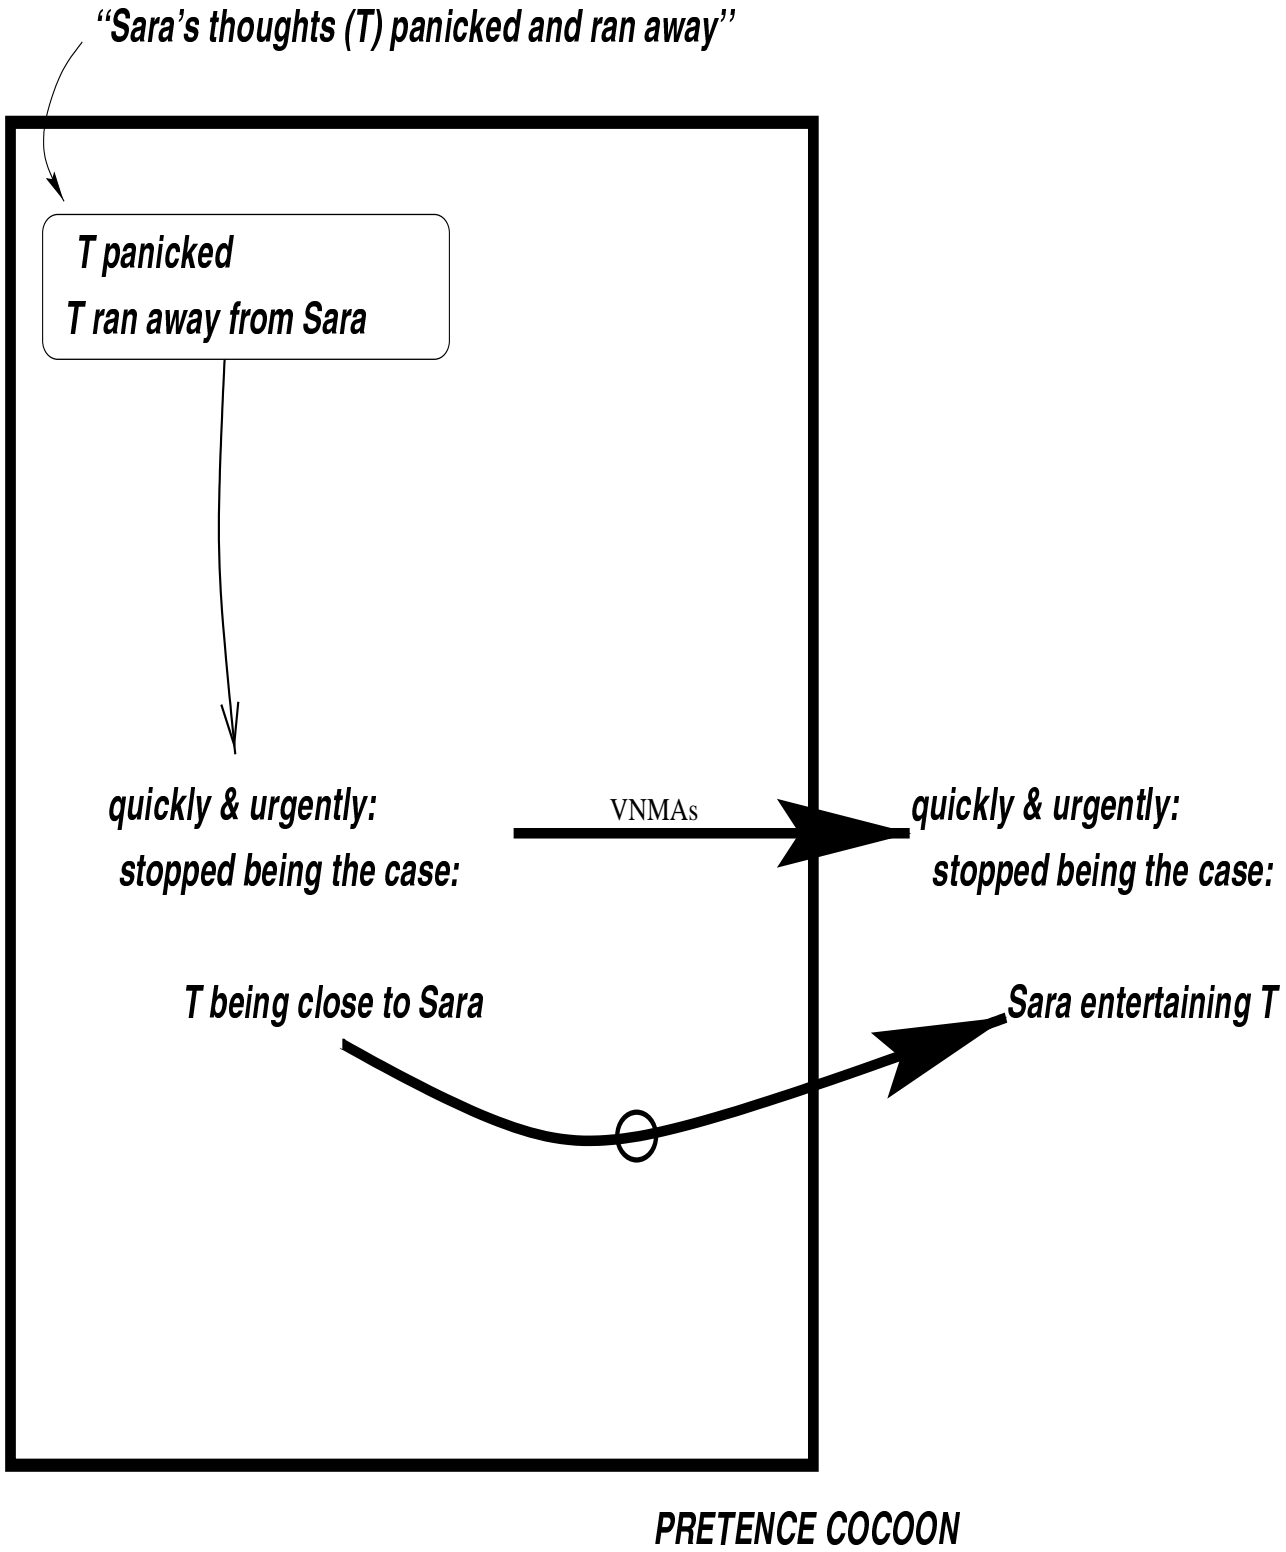


Figure 3: **Ideas as Persons or Other Animate Beings.** *The Panicky Thoughts.*

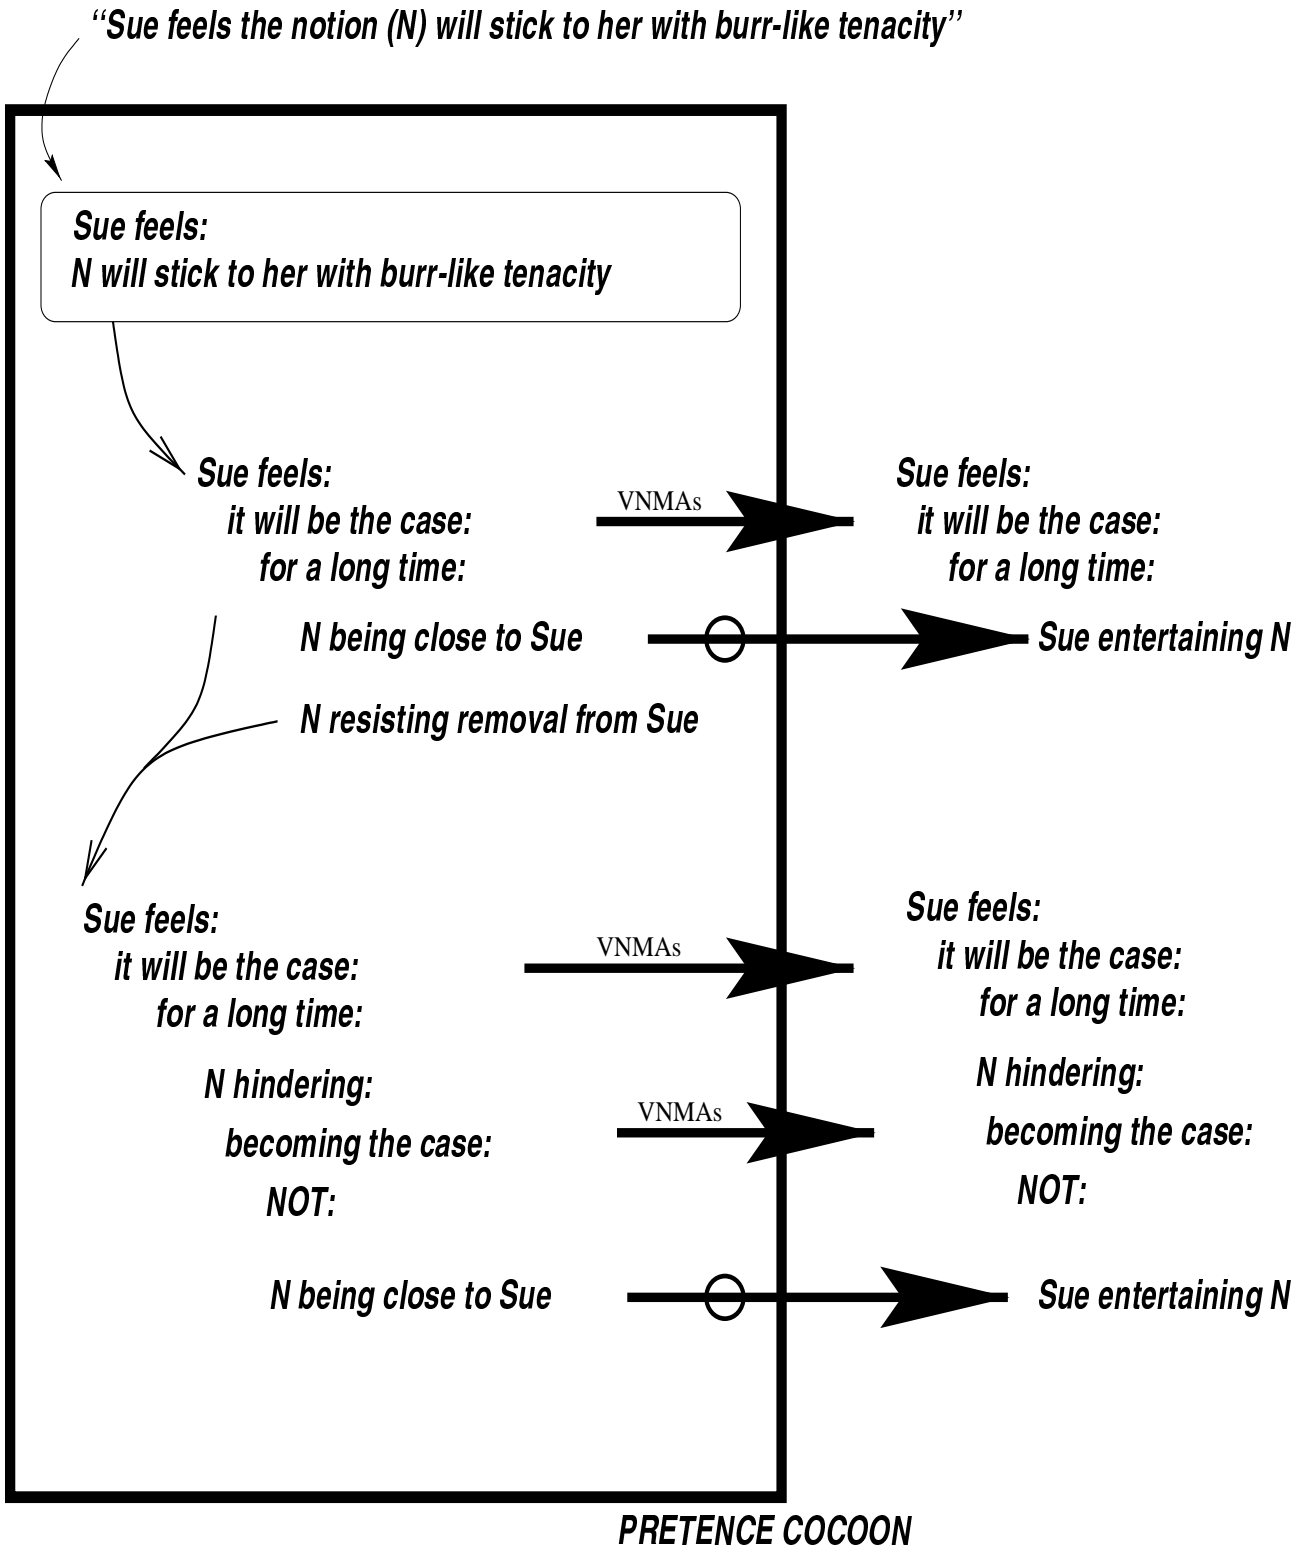


Figure 4: **Ideas as External Entities.** *The Burr-Like Notion.*

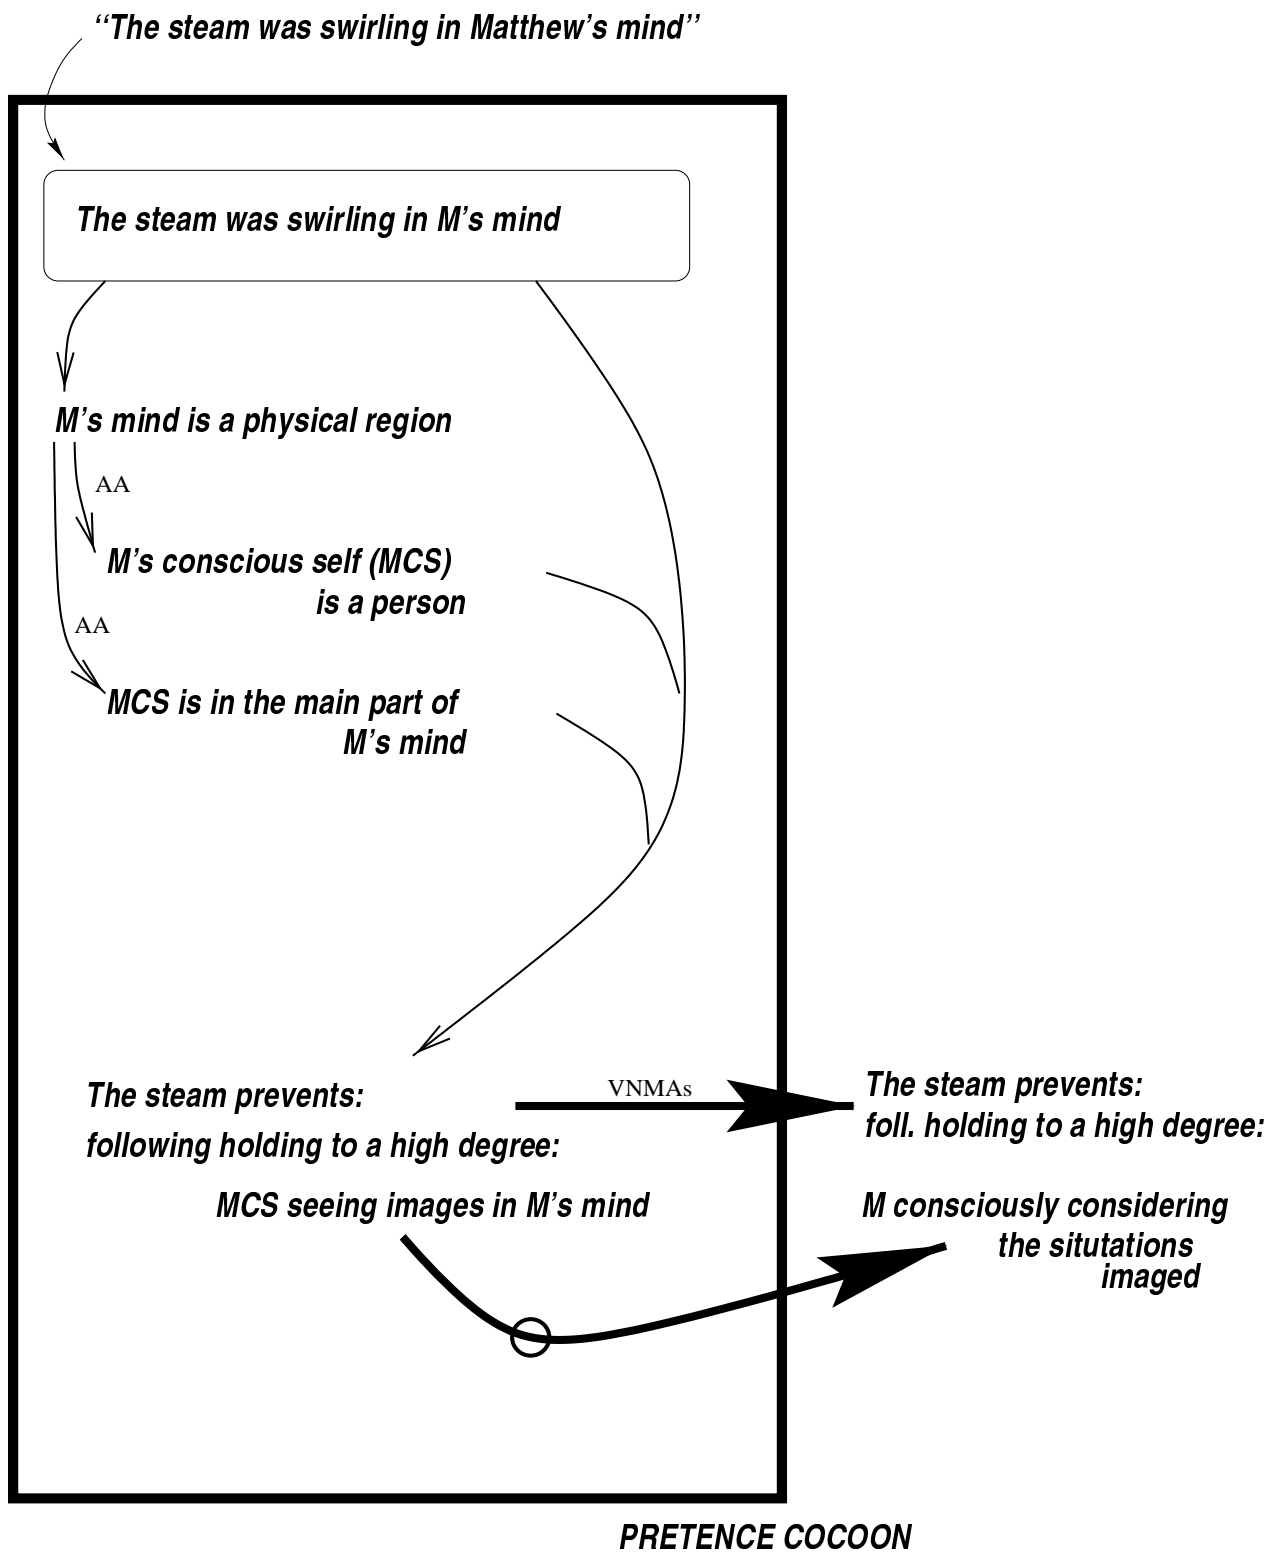


Figure 5: **Cognizing as Seeing**. *Steamy Thoughts*. The arrows marked AA depict applications of “ancillary assumptions” that form part of the metaphorical view of MIND AS PHYSICAL SPACE. Such assumptions are further discussed in Barnden & Lee (2001a).

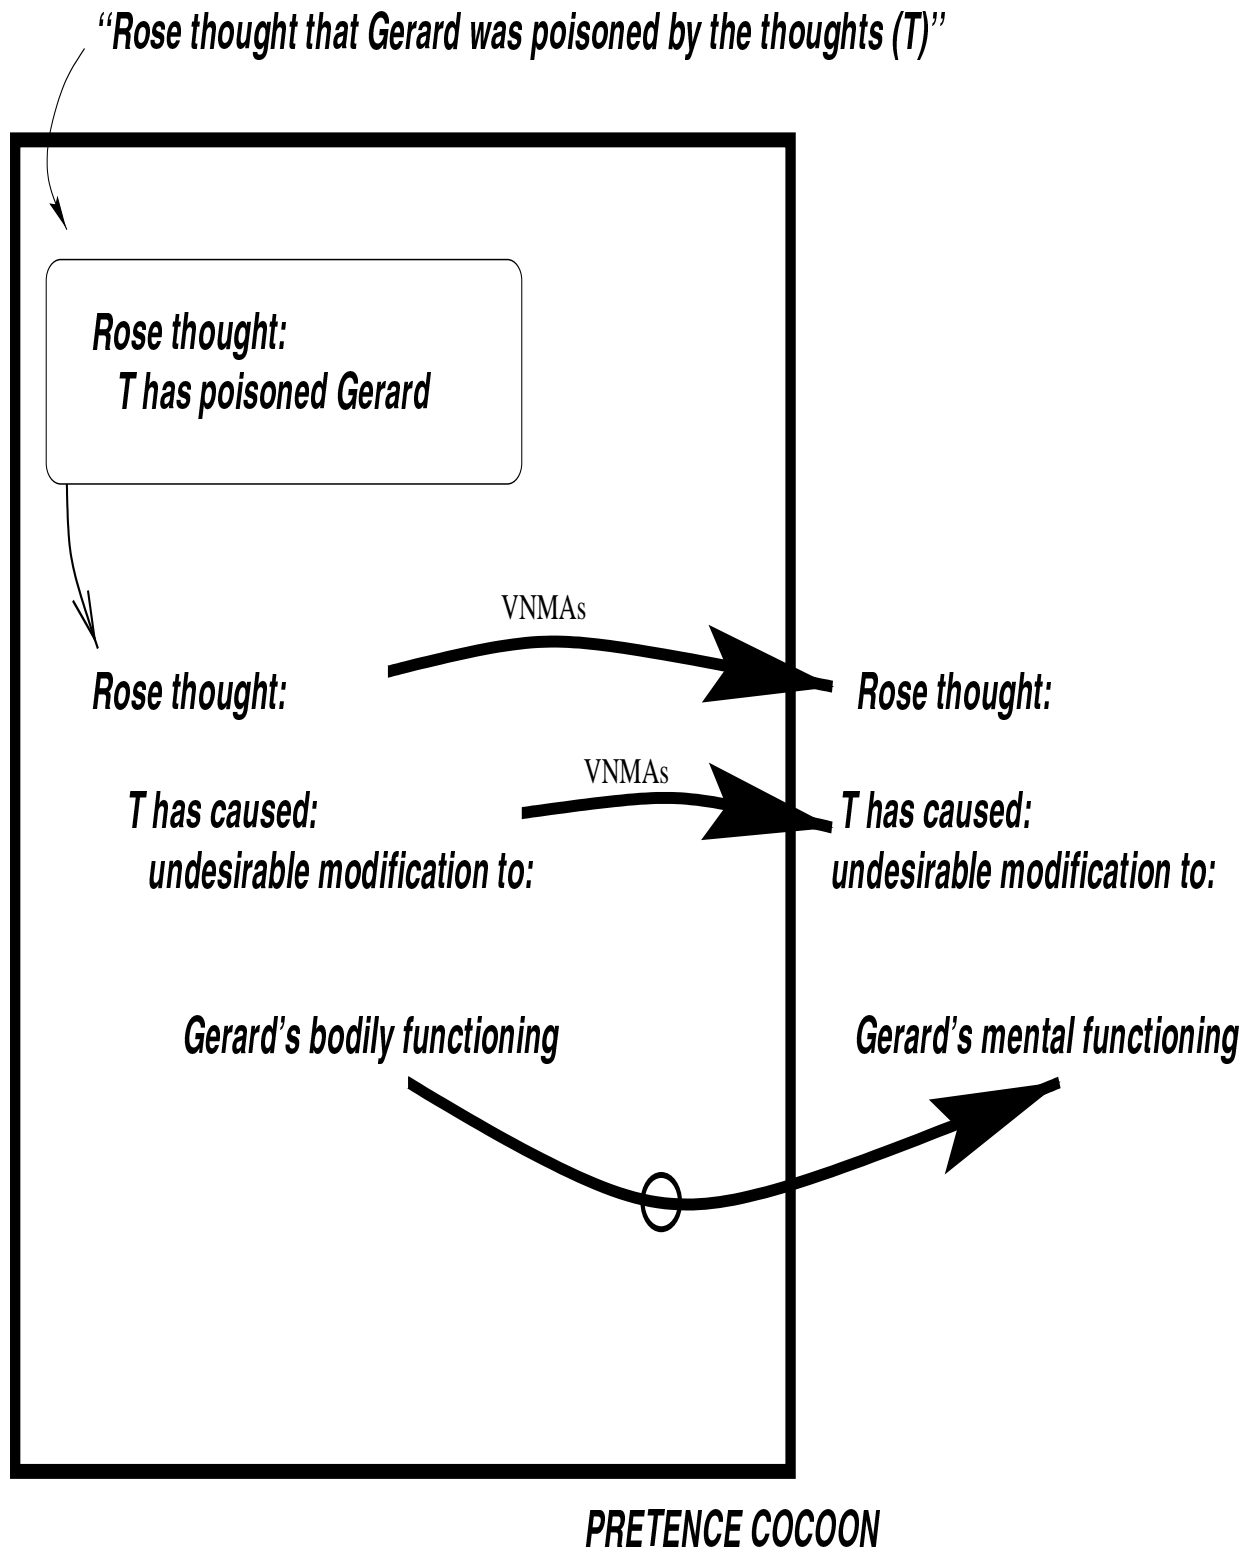


Figure 6: **Mind as Animate Being or Living Body. Mind Poison.**

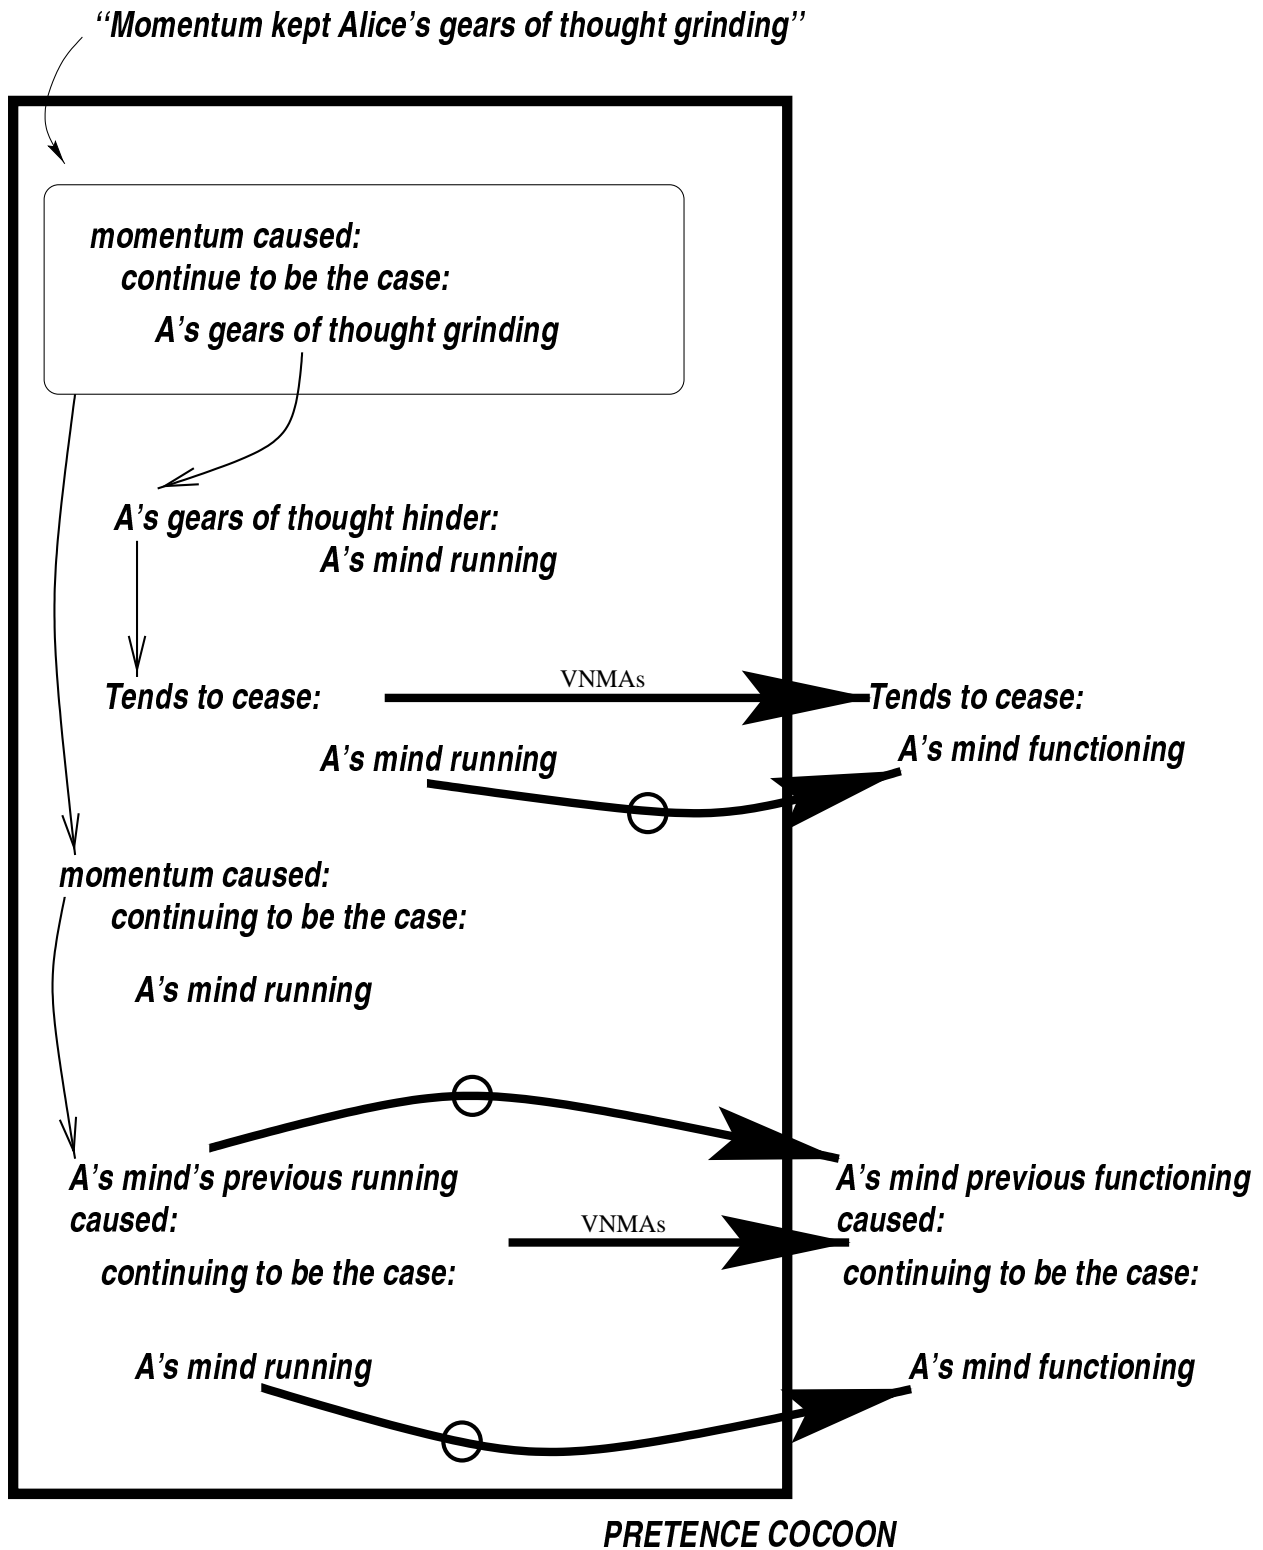


Figure 7: Mind as Physical Object—Other Cases: Mind as Machine. *Gears of Thought.*